IOmark-VDI

Tintri T5060 Storage System

Test Report: VDI-160826-a

Test Report Date: 26, August 2016

Table of Contents

Executive Summary 3
Vendor Supplied Product Description 3
IOmark-VDI Test Summary 4
IOmark-VDI Results 5
Tested Configuration Details 6
Hypervisor Configuration for IOmark-VDI Workload 6
Storage Configuration for IOmark-VDI Workload 7
Configuration Diagram 8
Connectivity 9
Tested Configuration Pricing 10
Detailed Results 11
Why the Need for IOmark-VDI 12
How IOmark-VDI Operates 12
Benchmark Application Workload Set 13
VDI Workload 13
Understanding Results 14
Benchmark Criteria 14
More Information about IOmark-VDI 14
About Evaluator Group 14

Executive Summary

This document is the official benchmark report for the tested configuration with a Tintri T5060 all-flash storage system running the IOmark-VDI benchmark.

IOmark is a storage specific workload and benchmark designed to test storage systems performance using a variety of real world, application centric workloads. The IOmark-VDI benchmark is a specific workload, which measures virtual desktop workloads (VDI) run against storage systems. Results are published after audit and certified approval by IOmark authorized auditors.

IOmark-VDI is a benchmark that certifies storage systems for virtual desktop workloads. The measurement criterion is storage performance, with the restriction that all storage workloads must be supported by the tested system. Although there are CPU and memory considerations, these aspects are not tested by the IOmark-VDI workload.

The results achieved by the Tintri system running IOmark-VDI are as follows:

- The T5060 supported 3,000 IOmark-VDI Office at a cost of $75.00 per VDI instance

A full description of the configurations tested along with pricing information is provided in this document. The criteria and performance requirements are as follows:

- For all application workloads:
  - All workloads must reside entirely on the tested system
  - 70% of response times for I/O's must not exceed 30ms
  - The replay time must complete within 1 hour and 15 seconds for each 1-hour workload

Vendor Supplied Product Description

Tintri simply stores virtualized workloads— offering a fully integrated VM-Aware Storage (VAS) system for virtualized enterprises and cloud. Balance your projects and workloads across high-performance All-Flash and award-winning Hybrid-Flash. Encrypt everything with a single click. Replicate individual VMs to remote locations. Sync child VMs with master VMs to speed development cycles. And manage it ALL from a single pane of glass. That's how you keep storage simple.

System design, including:

- Tintri VMstore UI, drill down to VM and virtual disk, in Tintri web-console or Hypervisor plugin
- Support for multiple, concurrent Hypervisors (VMware, Hyper-V, KVM and XenServer)
- Hypervisor-embedded storage for ease of management and deployment using existing tools
- Consistent performance and latency with all-flash T5000 series
- Operations managed per VM, including QoS, snapshots, clones and replication

Tintri all-flash T5000 series systems provide the following enterprise storage features:

- Capacity from 6 TB to 92 TB raw, and 17 - 308 TB effective capacity with data reduction
- Tintri all-flash arrays assign I/O on a per VM basis
- Quality of service (QoS) settings on individual VMs to guarantee performance
- Data reduction includes deduplication, compression, thin provisioning, and zero copy optimization

IOmark-VDI Test Summary

For the tested configuration, the following data is provided

Item	Value
Testing Identifier:	VDI-160826-a
Product(s):	Tintri T5060
Test Sponsor:	Tintri
Auditor:	Evaluator Group Inc.

Table 1: Test Identifier Information

Item	Value
IOmark-VDI Version:	Version: IOmark-VDI 3.8
Testing Completed:	July 2016
Equipment Availability:	August 2015
Audit Certification Date:	26, August 2016
Report Date:	26, August 2016

Table 2: Test Revision and Dates

IOmark-VDI Results

Shown below are the IOmark-VDI results for the system under test. The definition and workload characteristics of the benchmark are provided in Appendix A.

Price information provided below is explained in detail in Table 8 in this report.

A VDI user may be configured to run in one of two modes:

- Fully Provisioned (No clones - Requires at least 14 GB of host provisioned capacity / VDI user before compression)
- Linked Clones (Hypervisor based clones - Requires 4 GB of host capacity / VDI user)
- **Note:** Actual capacity utilization on storage may be lower, depending on data reduction rates

For each configuration run, the results are reported. Tables 3 and 4 below show the number of users supported (as defined in Appendix A).

VDI Mode	IOmark-VDI Office Worker	Tested Logical Capacity	Total Price	Price / User
Linked Clone	N/A	N/A	N/A	N/A
Fully Provisioned	3,000	132.0 TB	$225,000.00	$75.00

Table 3: IOmark-VDI Office Price-Performance Results

The "Office" workload is measured during a steady-state period during the VDI workload. The measurement period for the "Office" workload discards the startup and end portions of the workload and measures a 60-minute window during the middle of a 180-minute workload. The response time summary is shown below in Table 4, with further details shown in Figures 3 and 4.

VDI Mode	Test RAID Level	Average Read Resp. Time	Average Write Resp. Time	Average Response Time / VDI User
Linked Clone	N/A	N/A	N/A	N/A
Fully Provisioned	RAID 6	2.33 ms	2.23 ms	2.41 ms

Table 4: IOmark-VDI Office Result Details

As shown above in Table 4, the 3,000 Fully Provisioned Office users had an average response time of 2.41 ms, which is lower (better) than the required response time.

Tested Configuration Details

This section covers the connectivity, configuration and pricing information for the system under test.

Hypervisor Configuration for IOmark-VDI Workload

- A single Tintri T5060 system was used
- Multiple mount points were created to the two hosts used for testing
- A virtual disk was created for each of the reported IOmark-VDI's certified (3000 total)
- Data reduction was enabled
    - Utilized Tintri clones for data deduplication rates of approximately 10:1
    - Data generated was 2:1 compressible
    - Thin provisioning was enabled, saving additional capacity
    - RAID level utilized is not user selectable, used Tintri default settings

Detailed configuration parameters for the system under test, including connectivity are provided below in Table 5.

Storage System Parameter	Value
Hypervisor	VMware vSphere™ ESXi 6
Number of interfaces to the storage system:	4 Per Node (2 / controller = 4 total)
Connectivity to storage system:	2 @ 10Gb Ethernet / controller
Hypervisor storage protocol used:	NFSv3
Hypervisor version:	VMware ESXi 6.0U2 (3620759)
Thin provisioning:	Utilized in Tintri datastore
Hypervisor Storage Access:	NFS datastore
Datastore Filesystem:	NFS - filesystem access to Tintri
VAAI:	Yes, using Tintri VAAI plugin drivers
SATP:	N/A
PSP:	N/A
Total capacity of system allocated to IOmark-VM:	6.6 TiB usable (7.2 TB)

Table 5: Hypervisor Configuration Parameters

NOTE: Per IOmark requirements, a "write-only" workload is run prior to the actual workload. This pre-writes data to all storage locations referenced during testing. By pre-writing data prior to actual workload testing, there is no write allocation penalty associated with thin provisioning. This also ensures that when reads are performed the storage system reads the media, rather than returning zero's for unallocated addresses.

Storage Configuration for IOmark-VDI Workload

- A single Tintri T5060 provided the pooled capacity across both of the host test nodes
- 6 datastores were created on each host for testing (12 total)
- Each VM's VMDK was allocated using "thin provisioning" per Tintri and NFS default

Detailed Storage System configuration parameters for the storage system under test, including connectivity is provided below in Table 6.

Storage System Parameter	Value
Storage System firmware	Tintri 4.2.0.6
High Availability Access	Yes (active / standby)
Total raw capacity of system under test (SUT)	12 TB
Total usable capacity of system under test (SUT)	6.6 TiB (7.2 TB) usable, up to 35+ TB effective with data reduction
Datastores	Total of 6 datastores / host (No LUNs required)
Thin provisioning:	Yes
RAID Level(s)	Tintri default, RAID 6
Total Cache Capacity:	N/A
Read Cache:	N/A
Write Cache:	N/A
VAAI Features Enabled:	Yes
- NFS Full Clone	Yes
- NFS Extended Stats	Yes
- NFS Reserve Space	Yes
Automated tiering within the storage system:	N/A (T5060 is all-flash)
Deduplication or compression of data:	Yes, both in-line and always on
Storage system clones / writeable snapshots:	Yes, utilized storage clones during testing
Type of storage system clone:	Tintri native clones
Storage Media Utilized:	-
- SSD's	24 x 480 GB (includes spares)
- 15K RPM	NA
- 10K RPM	NA
- 7.2K RPM	NA

Table 6: Storage System Configuration Parameters

Configuration Diagram

The logical data layout of the test configuration is shown below in Figure 1. Since Tintri is an NFS datastore, block LUNs or volumes were not utilized. Instead, individual virtual disks were assigned to VM's as required, distributed across the 12 logical mount points. The VMDK's for each VM's workload were all allocated from the same Tintri Datastore capacity pool created by Tintri across the compute nodes.

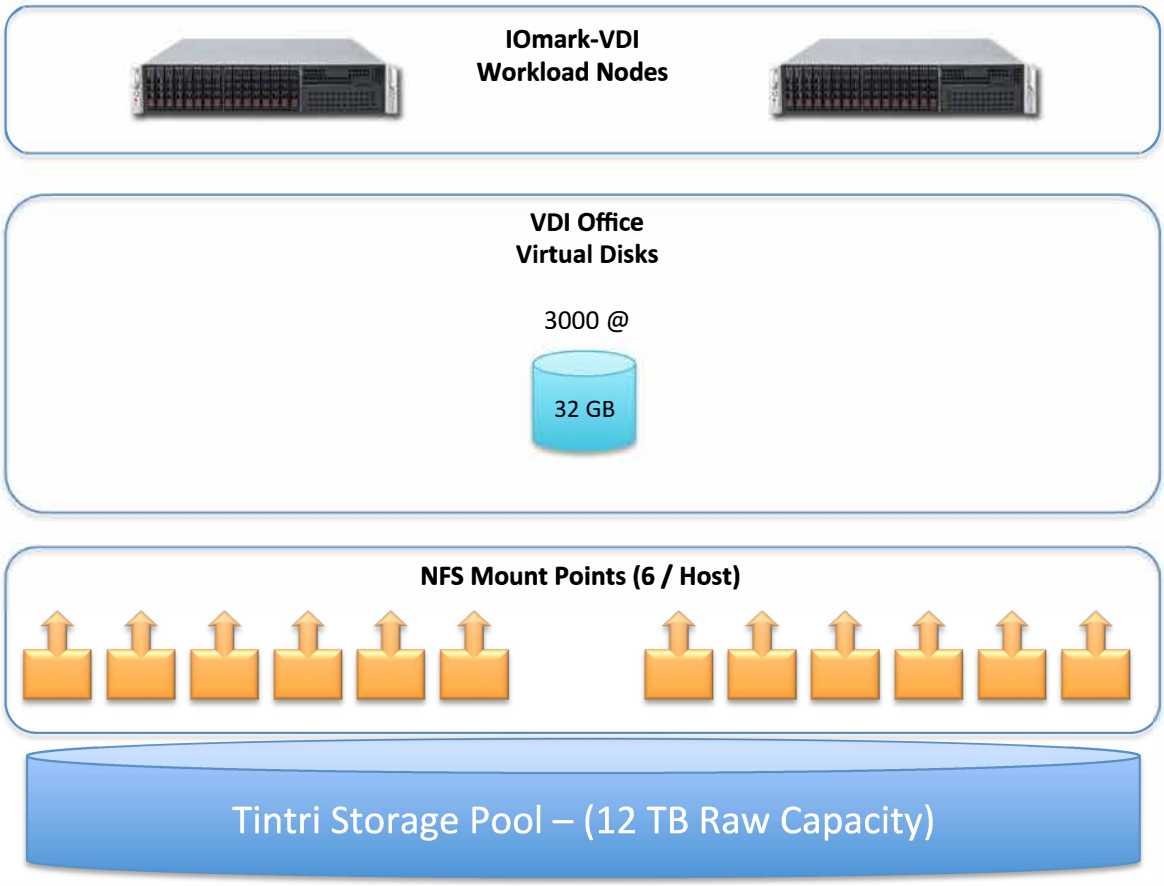

Figure 1: Logical System Configuration

Note: *Although a single NFS datastore could have been used, it was determined that utilizing more mount points improved performance. A total of 12 mount points were used, which equates to 12 host systems mounting a single Tintri shared datastore.*

Connectivity

Storage connectivity used was 10GbE, using LACP to bond the two network interfaces per controller. Each test host used 2 10GbE links to a 10 Gb Ethernet switch, for a total of 8 connections across both the two hosts and the dual controller T5060. Testing did not utilize a redundant HA configuration, although production deployment assumes connectivity to a HA network infrastructure.

The tested configuration connectivity diagram is shown below in Figure 2.

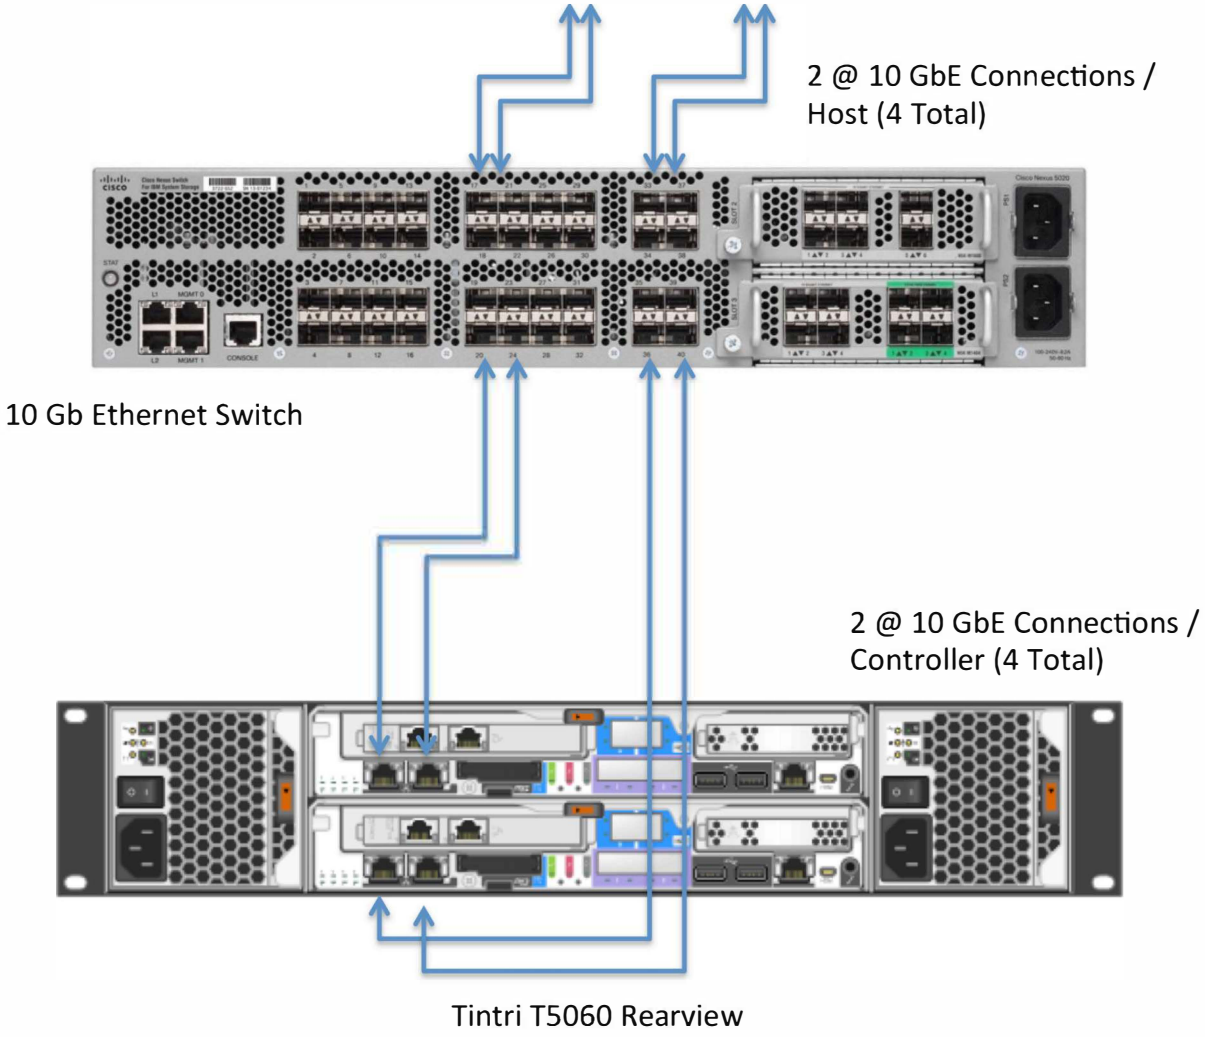

Figure 2: Physical System Connectivity

Tested Configuration Pricing

Item	Description	Qty.	Ext. List Price
1	Tintri T5060 (Includes base software)	1 system	$225,000.00
2	Tintri 3 years maintenance	Included	N/A
Total	**List Price HW + SW + 3 year service & support**		**$225,000.00**

Table 7: IOmark-VDI Price Information for Tested Tintri T5060 Configuration

Note: Support included for all hardware and software

Detailed Results

IOmark-VDI performance results are measured by application workload. The cumulative response times of all 3,000 Office VDI users that comprised the test workload are shown below in Figure 3.

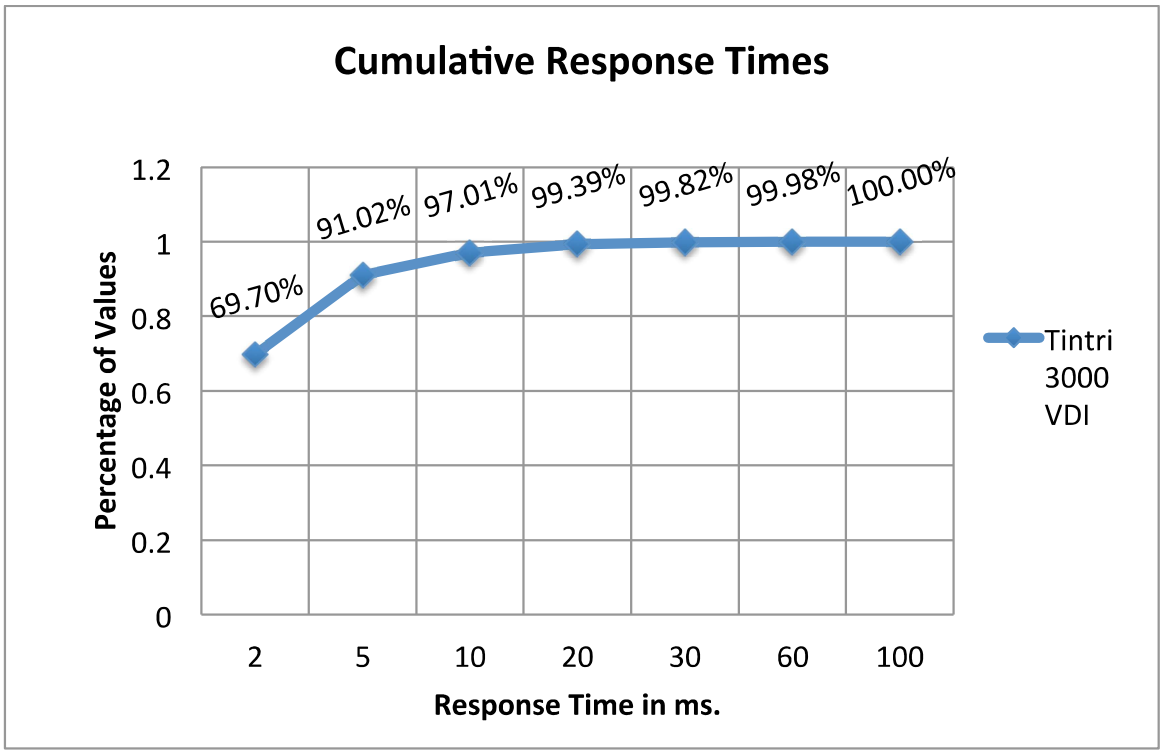

Figure 3: Percentage of Total Response Times at Measured Value

From Figure 3 above, the primary response time(s) of interest are:

- More than 90% of response times were less than 5 ms.
- 99.82% of response times were less than 30 ms. (exceeding requirements of 70%)

Appendix A - IOmark-VDI Overview

The ability to recreate a known workload is important for comparing a system against potential alternatives. Establishing a reference or benchmark workload enables system vendors as well as resellers and IT users to compare several systems utilizing a known workload.

Specifically, the IOmark-VDI benchmark recreates a storage workload that typically occurs in virtual desktop infrastructure environments. The workload is non-synthetic and recreates several applications that are commonly found in virtualized server environments.

Why the Need for IOmark-VDI

Currently, several application generators have been developed that are able to generate VDI workloads. However, there is no standard reference configuration, with the primary focus is on the server infrastructure. There are no existing benchmark workloads focusing on storage and storage system performance while running VDI applications.

By establishing a set of standard applications and capturing their I/O streams, it is possible to recreate application based storage workloads for these complex environments. IOmark-VDI is designed utilizing these concepts, and as such is the first benchmark designed to accurately generate application workloads for storage systems, enabling direct comparison of storage system configurations and their ability to support a specific number of applications.

How IOmark-VDI Operates

IOmark-VDI uses the concept of workload replay. I/O streams are captured from actual running applications and then "replayed" so that the exact sequence and I/O commands are issued. This allows the creation of a workload that is indistinguishable from an actual workload to the system under test, while being reproducible and requiring fewer resources. Additionally, the test environment is less expensive, easier and faster to create since actual applications are not required. Because CPU and memory are not consumed running applications, a much higher I/O workload may be generated with a set of server resources than is possible using native applications. This ratio is typically 10:1, but may vary.

In order to scale up the workload on a storage system, additional VDI workloads may be added to the same, or to other physical hosts. The only limitation to the scale of the test is the physical infrastructure supporting the workload. Sufficient, CPU, memory and I/O capabilities must be available to run additional workload sets.

Unlike artificial workload generation tools, IOmark-VDI recreates accurate read vs. write and random vs. sequential I/O requests. Another benefit of IOmark-VDI is the fact that it creates accurate access patterns, thus enabling storage cache algorithms to work properly.

Finally, IOmark-VDI maintains an accurate ratio of performance to capacity as workloads are scaled, ensuring that storage performance is measured with respect to storage capacity accurately. As a result, IOmark-VDI maintains an accurate ratio of I/O to capacity, producing results applicable to IT users.

Benchmark Application Workload Set

VDI Workload

1. View steady state operation
   a. Heavy Worker Profile – Average / VDI User
      i. 12.52 iops. / User
      ii. 1.06 MBps / User
   b. Standard Worker Profile – Average / VDI User
      i. 6.26 iops. / User
      ii. 0.53 MBps / User
2. Benchmark Criteria:
   - 70% of I/O response times must not exceed 30ms
   - All storage utilized must reside on/within the storage system under test

VDI Benchmark Parameters

- Operating System disk size is 20 GB (thinly provisioned)
- All user sessions were running Windows 7 as their guest OS
- No user data disk utilized
- VMware Linked clones may be utilized (as noted)
- Storage linked clones may be utilized (as noted)
- Heavy Profile:
  - The workload is non synthetic, actual I/O patterns are issued based on application capture
  - The size of I/O's is variable, ranging from 512, up to 2 MB transfers based on application
- Standard User Profile:
  - The workload is non synthetic, actual I/O patterns are issued as captured
  - Rates are 50% of "Heavy" user profile
  - The size of I/O's is variable, ranging from 512, up to 2 MB transfers

VDI Workload Generation

The workload generator used to generate the VDI workload was VMware View Planner. This application workload generator controlled running the 8 listed applications above, in a Windows 7 64 bit OS environment, running as a guest VM in a hypervisor environment.

VDI Workload Details

The specific applications comprising a VDI workload set are detailed below in Table 8.

Application	Storage Capacity / Instance
Guest OS (Microsoft Win 7 64bit)	20 GB
MS Office (Word, Excel, PowerPoint and Outlook)	N/A
MS Internet Explorer	N/A
Adobe Acrobat Reader	N/A
Windows Media Server	N/A
Windows 7 zip	N/A
Total VDI Guest Environment	Total = 20 GB

Table 8: IOmark-VDI Guest Application Overview

The total capacity required for each VDI user is approximately 20 GB of logical capacity. The capacity required for linked clone users is 4 GB, and the capacity required for fully provisioned users is 14 GB without data deduplication or compression.

Understanding Results

IOmark-VDI produces results indicating the response time of a storage system given a particular workload. Based on established criteria, these results indicate how many VDI sessions are supported by a specific storage configuration with a maximum allowed response time. The report is audited for accuracy and issued by Evaluator Group, Inc., an independent storage analyst firm.

Benchmark Criteria

IOmark has established the benchmark criteria for the IOmark-VDI workload. The performance requirements are established as follows:

- For all application workloads:
  - 70% of response times for I/O's must not exceed 30ms
  - All storage must reside on the storage system under test
  - The replay time must complete within 1 hour and 15 seconds for each 1 hour workload

More Information about IOmark-VDI

For more information about the IOmark benchmark, a theory of operations guide, published results and more, visit the official website at http://www.iomark.org. Some content is restricted to registered users, so please register on the site to obtain all available information and the latest results.

About Evaluator Group

Evaluator Group Inc. is a technology research and advisory company covering Information Management, Storage and Systems. Executives and IT Managers use us daily to make informed decisions to architect and purchase systems supporting their digital data. We get beyond the technology landscape by defining requirements and knowing the products in-depth along with the intricacies that dictate long-term successful strategies. Web: www.evaluatorgroup.com Twitter: @evaluator_group